Sketch-Sketch Revolution: An Engaging Tutorial System for Guided Sketching and Application Learning

Jennifer Fernquist[1], *Tovi Grossman*[2], *George Fitzmaurice*[2]

Autodesk Research
[1]jen.fernquist@gmail.com, [2]{firstname.lastname}@autodesk.com

ABSTRACT

We describe Sketch-Sketch Revolution, a new tutorial system that allows any user to experience the success of drawing content previously created by an expert artist. Sketch-Sketch Revolution not only guides users through the application user interface, it also provides assistance with the actual sketching. In addition, the system offers an authoring tool that enables artists to create content and then automatically generates a tutorial from their recorded workflow history. Sketch-Sketch Revolution is a unique hybrid tutorial system that combines in-product, content-centric and reactive tutorial methods to provide an engaging learning experience. A qualitative user study showed that our system successfully taught users how to interact with a drawing application user interface, gave users confidence they could recreate expert content, and was uniformly considered useful and easy to use.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors, Algorithms.

Keywords: Tutorials, Learning, Sketching, Engagement.

INTRODUCTION

Many software applications are difficult and intimidating for new users to learn. As a result, a significant focus of HCI research has been to not only improve the usability of software systems, but also their learnability [15]. One of the most traditional means to addressing software learnability is to provide application tutorials, which can guide users through tasks to obtain specific skills and experience with the system [13, 16, 19, 21].

Tutorials for design software, such as image editing or animation software, are shared among community members to showcase work and associated workflows. These forms of tutorials, which we define as *content-centric tutorials*, seem particularly promising in terms of engaging users. Their prevalence on community websites and discussion boards indicates that there is something compelling about learning by following along with someone else's workflow.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
*UIST '11,*October 16-19, 2011, Santa Barbara, CA, USA.
Copyright © 2011 ACM 978-1-4503-0716-1/11/10...$10.00.

Tutorials for *drawing applications* are particularly interesting to consider because the ability to complete a tutorial not only requires skill with the software *UI*, but may require *domain* artistic skills such as sketching. While sketching is something almost everyone wishes they could do well, the majority of people are not good at it [4]. This presents a challenge for tutorial design for drawing applications—if the challenges presented within these tutorials do not have a good balance with the user's skill level, then the user may not be fully engaged [5, 8, 9]. In particular, it may be difficult for a novice user to get a sense of the holistic, high-level drawing process and to "feel the experience" because they encounter difficulties at individual steps or because it takes an inordinate amount of time to complete a drawing. A more enjoyable experience may be possible if the tutorial could simulate for a new user what a work session would be like if they were already a skilled artist.

In this paper, we present Sketch-Sketch Revolution, an engaging, content-centric tutorial system for drawing applications that allows a user, of any skill level, to experience the success of creating high-quality drawings by providing an adaptable level of *UI*, *domain*, and *content* assistance. The system is built onto SketchBook Pro, an existing application for creating sketches and drawings using digital tools that mimic traditional media. Sketch-Sketch Revolution gives users the opportunity to learn about the application's user interface and its associated workflows and also get exposure to drawing techniques and related foundational concepts. A user study showed that our system taught users how to interact with a drawing application user interface, gave users confidence they could recreate expert content, and was uniformly considered useful and easy to use.

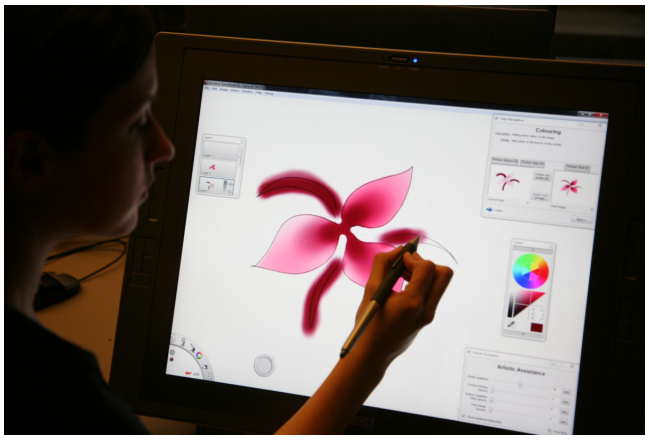

Figure 1. Working with Sketch-Sketch Revolution.

RELATED WORK

Tutorial Systems

Early HCI research quickly recognized the challenge in providing effective help systems for software applications [7]. Help systems quickly evolved, taking into account research proposing minimalist and task-centered documentation [6, 7,28].

More recently, researchers have argued that help systems should be provided within closer context of the actual application [16, 21]. For example, Tooltips [10], and more recently ToolClips [16], provide pop-up help within the application. However, the help which they provide is feature-centric, whereas in our work we hope to provide tutorials for entire high-level tasks.

The Graphstracts help system [19] provides screenshot snippets of the application to help users understand a task sequence. However, Graphstracts were also designed for low-level tasks and were provided outside of the actual application. Some systems, such as Document Wizards provide in-application assistance, but typically this is limited to very simple tasks [27].

Possibly most closely related to our work are the Stencils-Based Tutorials [21]. These tutorials are presented within an application and help users complete an entire task from beginning to end. Furthermore, they provide assistance by forcing a user to perform the correct behavior by overlaying a stencil with a hole to click-through. Google Sketchup's self-paced tutorials [13] also provide a similar experience, but with weaker guidance. These examples serve as important inspiration and, in our work, we explore how such tutorials can be extended to handle more advanced interactions, particularly sketching. We will also explore other new areas of the design space not previously investigated, such as how such tutorials can be configured to match a user's expertise level, how higher level information about the application or domain can be integrated into the tutorials, and how to handle areas of a tutorial that may become redundant or boring for users.

Sketching/Art Assistance

Relevant to our work are systems that have been designed to help users do artistic work. There has been a long line of research on sketch beautification. Examples include the Pegasus [20] and Fluid Sketches [2] systems, which take a user's freehand stroke and beautify it by considering geometric constraints [20]. There are also systems which use predefined geometric templates to aid sketching, such as Digital French Curves [26] and Kinematic templates [12]. We will be presenting a new algorithm to beautify strokes using an author's freehand stroke as a template.

There have also been systems to help users create artistic content. The ILoveSketch [3] and EverybodyLovesSketch [4] systems provide an interface to allow users to create 3D sketches from a 2D perspective viewpoint. Also relevant to our own work are the projector-guided painting [11] and ShadowDraw projects [22]. With these, users are guided by background imagery to recreate an author's artistic content. However, in both systems authors are purely sketching on top of the author's content and are not seeing or experiencing how the author actually created the artwork. In our system, users will be able to experience exactly how authors created content, such as the order and temporal dynamics of their strokes, and how they leveraged software features to obtain their final content.

Capturing and Sharing Workflows

There have been a number of projects that have attempted to automate the process of tutorial creation to mitigate the time required to create tutorial content. Automated Photo Manipulation Tutorials records a user's workflow and automatically creates a static text and image based tutorial [14]. While the work was limited to short tutorials ranging from 4-6 steps, the concept is inspiring and we hope to extend it into full workflow tutorials. In contrast, Chronicle tracks the entire workflow history of a document, including a video screen capture, but does not provide the resulting content as a step-by-step tutorial [17]. Other related projects have looked at recording, visualizing, and navigating software workflows, but not with the intention of using those recordings to create interactive tutorials [23, 24]. Stencils-based tutorials did provide a tool to create tutorials, but authors had to manually enter information about each step. Our work extends these systems by combining the benefits of the automatic tutorial creation of photo manipulation tutorials, the full workflow capture of Chronicle, and the in-application experience of Stencils-based tutorials. Furthermore, our system will assist its users when they try to recreate the actual captured content, in addition to helping with the user interface steps.

Creating Optimal Experiences

One of the goals of our system will be to create a compelling experience so users stay engaged. In psychology, the concept of Flow has been used to describe a state of "optimal experience" [9] (Figure 2). The concept of flow is commonly used to guide HCI research and design [5]. However, we are unaware of research which applies these concepts to the *learning material* presented for user interfaces. Chen et al. [8] identify 3 qualifying factors of the activity itself for reaching the flow state*: perceptions of clear goals*, *immediate feedback*, and *matched skills and challenges*. By designing a tutorial experience that provides these 3 characteristics, users will have the opportunity to achieve an intrinsically rewarding experience [9].

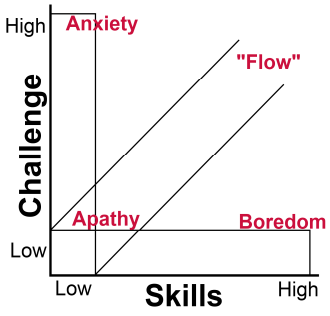

Figure 2. Graphical depiction of Flow [9].

DESIGN SPACE OF INTERACTIVE TUTORIALS

As discussed in our review of related work, there is a wide spectrum of tutorial formats that can be presented to a user. The design space presented below is not meant to provide a full taxonomy of tutorials, but is meant to serve as a summary of the dimensions most relevant to our work and highlight the properties we will be utilizing.

Scope

At the lowest level, tutorials can focus on an individual tool or feature of a system (*feature centric*) [16]. They can also demonstrate how a low-level task can be accomplished (*task-centric*) [14]. At the highest level, tutorials can describe how an entire workflow is used, from beginning to end, to create content (*content-centric*) [21]. We will be focusing on ***content-centric tutorials***.

Interactivity

Traditional tutorials are *passive* if the user only consumes the learning material, such as reading an article or viewing a video [14]. Alternatively, we define tutorials that provide a mechanism to allow users to try out the concepts they are learning as *active* [13]. Lastly, *active* tutorials that are aware of a user's interactions and can respond are considered *reactive* [21]. Our focus will be on ***reactive tutorials***.

Integration

Tutorial materials were initially provided through *physical* books and guides [7]. When the material is available on the computer system, the tutorials are often called *on-line* [19]. At the greatest level of integration, the tutorials can be provided *in-application* [27]. We will be focusing on ***in-application tutorials***.

DESIGN GOALS

In this section we outline the design goals we will adhere to in the design of our tutorials system. These goals are motivated by our high-level effort to create an enjoyable and rich learning experience.

D1. Highly Engaging. Our first goal is to design an experience that will be as compelling as possible. Building from our review of Flow, this will include giving users a perception of *clear goals* and *immediate feedback*.

D2. Provide an Adaptable Experience. Almost all tutorials we have reviewed are rigid; that is, their presentation is fixed regardless of the user's skill level. This reduces the likelihood that the challenge of the tutorial will match the user's skills. Our second design goal is thus to provide an adaptable experience so that the tutorial can be configured to provide a degree of challenge consistent with the user's skill level (Figure 2).

D3. Flexible Time Commitment. Tutorials for creating rewarding content may contain repetitive steps or steps that a user has already mastered. This goal is to provide a mechanism for users to dictate the amount of time they spend on the tutorial or any given step.

D4. Application Learning. A required goal of any application-focused tutorial system is to help users learn about the application. This should include learning where tools are located, how tools are used, and how multiple tools can be used together to perform workflows.

D5. Teach Users About Drawing. It is not the goal of our system to turn any user into a skilled artist. However, we do want to give users exposure to how real artists work. This will be a unique aspect of our system, as software tutorials typically focus on learning the application, not the associated domain knowledge.

D6. Enable Users to Experience Success. Tutorials may be frustrating and difficult to follow when a user lacks the skills required to succeed. In the sketching domain, we want to ensure success with the UI, and have an additional challenge of helping users with the content creation itself. Our final goal is to help users achieve success, to increase their enjoyment of the experience.

FOCUS GROUP MEETINGS

Although the goal for our system is not to provide intense artistic training, we do want to ensure that we adhere to any important fundamentals with respect to learning about drawing. As such, we held two focus group meetings with students and faculty from the Ontario College of Art and Design. During the first we presented the design concept for our system. Feedback received from this meeting confirmed some of the positive aspects of our system, but also resulted in changes to some specifications. The second meeting, in which they piloted a prototype, also elicited feedback resulting in minor modifications. Design decisions affected by these meetings will be highlighted as we describe the system below.

SKETCH-SKETCH REVOLUTION

In this section, we describe our new tutorial system, Sketch-Sketch Revolution.

Target Application and Implementation

SketchBook Pro[1] is a digital drawing application that has marking-menu interface elements and offers drawing tools such as pencils, brushes, and markers, and can be used with pressure-sensitive tablets and tablet PCs. Sketch-Sketch Revolution was implemented by making modifications to the SketchBook Pro source code.

System Overview

Sketch-Sketch Revolution offers two components: an authoring tool and a tutorial system. The authoring tool allows artists to create drawings in SketchBook Pro that are automatically converted into tutorial files, with minimal additional effort required on the author's part. The tutorial system loads an author tutorial file and automatically generates a custom tutorial experience. Users are guided through each step, informed when, where and how to change application settings, as well as when and where to draw, with feedback indicating their success at these tasks.

As we describe the features of the system, we indicate in parentheses the relevant design goals they address.

[1] http://www.autodesk.com/sketchbookpro

Authoring Experience

Similar to previous projects [14, 17], we wanted the authoring experience to be simple so artists could focus on drawing. The only manual input *required* from the author is assigning a tutorial name. All information described below is exported to a custom XML file when capturing is complete.

Capturing Workflows

SketchBook Pro has an existing journaling capability to store all user operations. Sketch-Sketch Revolution utilizes this to automatically record the brush settings, layer manipulations, and colour selections that the author carries out. It also records information for all of their brush strokes, including position, timing, and pressure data.

Author Descriptions

The author has the option to add descriptions to their workflows by entering text into the appropriate text fields in an Author Tool dialog. They can add a label, a description, and additional details. This gives authors the opportunity to describe supplementary information as to the motivation for the step, hints on how to proceed, or relevant domain knowledge (*D5*).

Foundational Concept Overlays

The tutorial author also has the option to choose a *foundational concept overlay* image from a preset list to attach to any step in the tutorial (*D5*). These overlays contain an example image with text briefly describing a fundamental drawing concept, such as lighting, 1- , 2-, or 3-point perspectives, or gauging proportions. We included 9 such concept images—5 created by professional artists specifically for our project, and 4 taken from art text books (Figure 3).

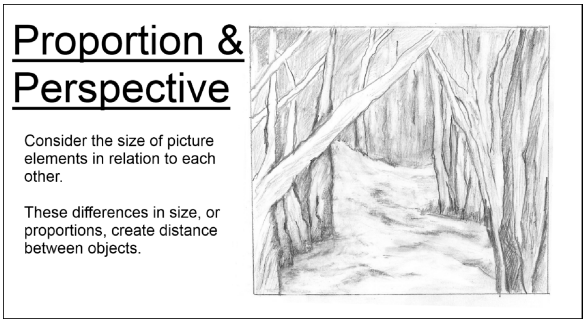
Proportion & Perspective
Consider the size of picture elements in relation to each other.
These differences in size, or proportions, create distance between objects.

Figure 3. One of nine foundational concept images that authors can associate with tutorial steps.

Tutorial Generation

A tutorial is started by opening an XML file created by the authoring tool. The system automatically parses the data and extracts the workflow information, author descriptions, and fundamental concepts. Regardless of the information authors provide, the tutorial is automatically segmented into steps. Each step consists of a UI phase and a drawing phase, defined by a consecutive set of UI actions followed by a consecutive set of drawing actions (0). Redundant actions (such as changing the brush size twice before drawing) are automatically removed. 'Undo' actions and the strokes they are undoing are also removed. While such actions may be part of the creative process [1], we omitted a mechanism for authors to choose to include undo's in the recorded workflow to simplify the experience for both authors and users.

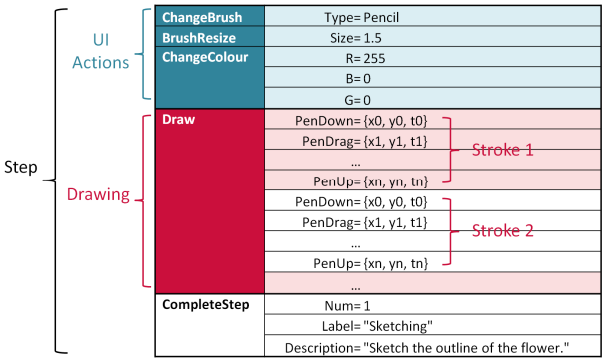

Figure 1. Pseudo code example of parsed XML. A step contains UI actions first, then a drawing action with 1+ strokes, and then a step completion entry.

Tutorial Navigation

Startup Splash Screen

Upon first opening a tutorial, the final author image is shown. Users begin the tutorial by pressing a Start button. The final author image disappears and users are presented with a blank canvas and a Step Navigation dialog.

Step Navigation Dialog

The Step Navigation dialog is used to lead the users through the tutorial (Figure 2). Any information the author added to the tutorial is displayed at the top of this dialog (*D5*). It updates automatically as the user progresses through the tutorial steps. Below, two thumbnail images are shown for reference: the completed step image, and the final image (*D1)*. The completed step image is automatically updated as the user progresses through the tutorial, displaying what the author's image looked like upon completion of the current step. At the bottom, the current step instruction is displayed. The Step Navigation dialog also contains several interactive buttons which will be described later.

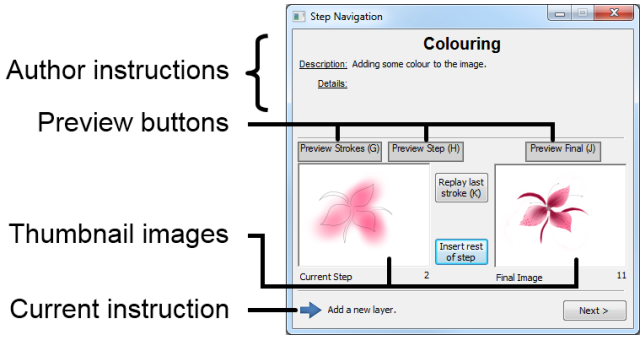

Figure 2. The Step Navigation dialog.

User Interface Phase

Positional Callout Instruction

Step instructions for user interface actions, such as applying appropriate brush settings or manipulating layers, are

also displayed in a separate callout dialog that is dynamically positioned next to the associated UI component. For actions involving complex interactions, such as selecting an item in a marking menu, informative images accompany the instruction (Figure 3). The tutorial system dynamically determines if the necessary dialog for a step is not open. If it is not, the callout instruction will first prompt the user to open it. These instructions help teach users where application elements are and how to use them (*D4*).

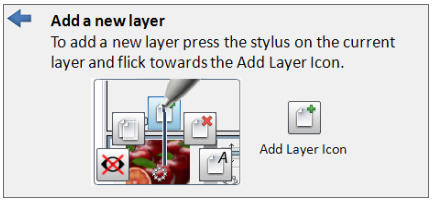

Figure 3. A positional callout instruction informing a user to add a new layer and guiding how.

When the user carries out the appropriate action, a checkmark appears on the screen to give positive feedback to the user and then the callout dialog text and position are updated to the next instruction (*D1 – immediate feedback*).

For steps that require setting a parameter to a continuous value, such as the brush size or colour, the user only needs to set the value within 5% of the target value for the tutorial to progress to the next step.

Sometimes, a user may want to continue to the next step without correctly completing the current one as instructed. This could be useful, for example, if they want to select a slightly different colour than the one the author used. To manually proceed, the user can click the *Next* button in the Step Navigation Dialog.

Drawing Phase

When all UI actions for a step have been carried out, the callout instruction dialog disappears and the Step Navigation instruction text says 'Draw!'

Preview Strokes

To proceed with a drawing action, the user presses the space bar. This causes a preview stroke to play, demonstrating the next stroke that the author drew in the original image. The stroke plays back at the same speed the author drew it. The stroke can be replayed as many times as desired by pressing the hotkey K or clicking the replay button in the Step Navigation dialog. The user can then attempt the stroke and proceed to the next via the space bar. Preview strokes are drawn on a separate layer so users only see authored content while previewing and to prevent synchronization issues between the user and author drawings.

When the user's brush type is a Pencil or Pen, the preview strokes are shown one at a time. When a Pencil/Pen stroke is finished previewing, its opacity decreases (default 50%) so that the user can partially see where they should draw next. The user can overdraw on their own stroke, which automatically removes their previous stroke, until they input a stroke they are happy with.

When the user's brush type is an Airbrush, Paintbrush, Chisel Tip Pen, or Eraser, then at least 3 seconds of preview strokes are shown, even if this spans multiple author strokes. We have found that for these brush types, some author strokes are so brief that it is difficult to discern tiny stylus flicks and too tedious to play them back individually. Opacity of these preview strokes is set to 0 when finished.

The *preview stroke* feature takes the sketching process beyond tracing. This is something we aimed to achieve based on feedback received from our focus group meetings. Tracing only allows users to see the final image they have to achieve. The preview strokes allow users to see, stroke by stroke, how the original image emerged, including the temporal dynamics that the original author used (*D5*). This is useful for users who are overwhelmed by the complexity of a drawing and do not know where to begin. We believe this may increase engagement (*D1*) by providing users with *clear goals*.

Stroke Guidance

One of the challenges with sketching tutorials is not mastering the UI steps, but the actual sketching. Sketch-Sketch Revolution provides *stroke guidance,* so that the content a user creates can resemble the artist's drawing (*D6*). When drawing with the Pen or Pencil, a user's strokes are modified to resemble the artist's original stroke.

We developed a new stroke correction algorithm to accomplish this. In SketchBook Pro, strokes are composed of a series of points and lines are interpolated between them. The algorithm consists of 3 steps. First, it analyses the points of the author's stroke and records the distances between consecutive points, relative to the entire length of the stroke. Second, the algorithm re-samples the user's stroke so that it contains the same number of vertices as the author's stroke, with the same relative distances between consecutive vertices. Third, a new stroke is generated, defined as linear interpolation between each pair of corresponding points of the author and user strokes. The extent of the interpolation is a single parameter in our algorithm that defines the strength of the guidance. When the guidance value is set to 0%, the original user stroke remains. When the guidance is set to 100%, the user's stroke will snap to the author's stroke. By default, our application sets the guidance value to 50% (Figure 4). The stroke guidance enables users to experience success in drawing.

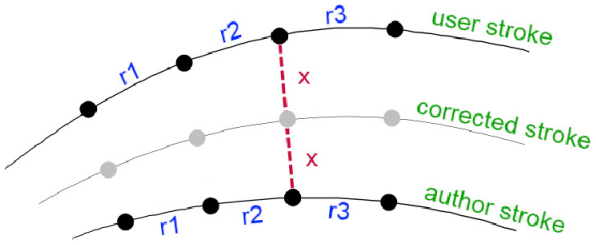

Figure 4. Example of correction applied to a user stroke when the *stroke guidance* value is 50%.

The stroke guidance algorithm resembles the $1 recognizer presented by Wobbrock et al. [29], but the target stroke is not resampled and a non-uniform resampling is used, which preserves the target stroke's original shape. While more advanced techniques could be considered, such as dynamic time warping, our technique is simple to implement and provides good results.

Stroke Feedback

In an effort to better engage users, we provide them with *immediate feedback* after each stroke is drawn (*D1*).

After each Pencil or Pen stroke the user draws, we calculate the average distance between corresponding points in the author's stroke and the re-sampled user stroke. This average pixel distance is subtracted from 100 and presented as a percentage score (e.g. a 90% match corresponds to an average pixel distance of 10). This score is presented to the user when each stroke is completed (Figure 5, left). The average score over the entire drawing is displayed to the user at the end of the tutorial.

We also provide timing feedback. The time difference is deemed to be the difference between the author and user stroke times divided by the author's stroke time. In our focus groups meetings, it was emphasized that drawing strokes faster or slower isn't necessarily an indication of correctness. Thus, we felt it may be useful to visualize the temporal comparison, but not to associate it with any sort of language or score. A range slider indicates the author stroke speed with a blue icon in the center (Figure 5, right). If the user is faster, a green icon appears to the right of center. If they're slower, the icon is to the left of center. The distance between the icons indicates speed difference between strokes, where each slider tick is 10%.

Figure 5. Stroke (left) and timing (right) feedback.

Automatic Step Completion

During a drawing phase, the user can choose to insert the author's strokes for the remainder of the step with the "Insert rest of step" button in the Step Navigation Dialog. The remaining author's strokes are automatically drawn into the current user's image layer. This feature facilitates a rapid walkthrough of a tutorial, allowing users to skip or quickly get through a step they find long and repetitive, without impacting the final content they generate (*D3, D6*).

Preview Images

During the tutorial, a user has the option to preview the author's content. They can preview what the author's content looked like up to the current stroke, to the end of the current step, and to the completion of the tutorial. The previews are overlaid on top of the user's canvas and are shown when the user either presses the preview buttons on the Step Navigation Dialog or hits their associated hotkeys. The previews are spring-loaded, disappearing as soon as the button or hotkey is released. These previews help provide users with *clear goals* (*D1)*.

Foundational Concepts

If an author has associated a foundational concept with a step, a button becomes visible in the Step Navigation dialog to indicate its availability. When clicked, the foundational concept image and text is overlaid on top of the user's image. It can be hidden by clicking the button again. This allows users to learn about important drawing concepts as they progresses through the tutorial (*D5*).

Adapting the Tutorial Experience

As set out by our design goal, the tutorial experience is adaptable (*D2*). Novice users can setup the tutorial to provide a heavier level of assistance while more advanced users can opt to work with minimal assistance using the Artistic Assistance Dialog.

The stroke guidance is adjustable via a slider, with the default value set to 50%. The opacities of the preview strokes, author step image, and author final step, are also adjustable via sliders. Users can continually view any combination of these images, at any transparency, to help guide them as they create their drawing.

Figure 6. A) The Artistic Assistance dialog and B) the Startup dialog.

The positional callout dialogs can be disabled via a checkbox in the Artistic Assistance dialog. This may be useful to users who are already familiar with the application UI, or to users who want an additional challenge.

A final option from this dialogue is to enter a *freeform mode* where users can draw or perform UI actions without impacting the tutorial. Strokes are not corrected or scored in this mode and steps do not automatically advance.

Users can also customize the amount of time that they want to dedicate to a tutorial using a dialogue shown at the startup splash screen. The user can either take the full, unmodified length, or enter a desired number of minutes in a text field (*D3*). This feature is useful if a user has a limited amount of time to complete a tutorial. When a custom time is set, the system divides the time equally amongst all steps. During the tutorial, once the system detects time has "run out" for a step, it automatically completes the remainder of the step and progresses to the next.

USER STUDY

We conducted a qualitative user study on Sketch-Sketch Revolution. Our goals were to obtain initial observations and feedback from users, understand which features they found most useful or easy to work with, identify any features that may require modification, and gauge users overall enjoyment experienced in drawing with our system.

We tested two versions of Sketch-Sketch Revolution. The first was a full version of our system, as described above. The second was an 'un-guided' version of the tutorial system that lacked stroke guidance, stroke feedback, and previewing. Positional callout instructions were provided, but without the integrated image-based instructions. We tested this version of the system to understand the impact of the most novel features in Sketch-Sketch Revolution in comparison to Stencils-based tutorials [21].

Participants

Eight right-handed participants (4 female) between the ages of 22 and 59 (μ=33.4) with normal colour vision were recruited via online and email postings. No participants had awareness of our project research goals. One participant had significant drawing experience, 3 had minimal experience, and 4 had none. Only one participant had used SketchBook Pro before.

Apparatus

Our user study was conducted with a Wacom Cintiq 21UX tablet, which has an active area of 43.2cm×32.4cm and a resolution of 1600×1200.

For the un-guided tutorial a secondary monitor, with a diagonal of 30" and resolution of 2560×1600, was used to display a 1600×1200 dialog that allowed users to see/toggle between the full sized preview step and final images. This was included to show users what to draw at each step.

Participants sat in front of the tablet, with a stylus in their dominant hand, a keyboard near their non-dominant hand, and the secondary monitor was to the left of the tablet.

Procedure and Design

The study was divided into two trials, lasting about 30 minutes each. In each trial, the user created one drawing with either the fully functioned Sketch-Sketch Revolution system (*Guided*) or the un-guided version (*Un-Guided*). Two sketches and their associated tutorials were created internally by authors with artistic experience – one of a fish, and one of an orchid. The order of the conditions, and the condition to sketch assignments were both fully counterbalanced. Users were briefly trained on the functionality of each system before beginning. Users answered a questionnaire after each trial, and a post-study questionnaire. Participants performed both trials in one session lasting approximately 60 minutes. An experimenter took observations and provided assistance when necessary.

Results and Observations

Questionnaire Results

The questionnaire at the end of each trial contained eight questions about the tutorial system for that trial, which participants were asked to rate on a 1-5 Likert scale (where higher is better) (Figure 7). Reactions to both systems were relatively high. The lowest response across all questions was for drawings being easy to complete, in the un-guided condition (μ=3.63). For guided tutorials, the response to this question was significantly higher (μ=4.5, $F(1,7)=6.24$, $p<.05$). This was the only significant difference in the two conditions, although as Figure 7 shows, responses for the guided system were consistently higher.

Also included on the questionnaire were six TLX-based Likert-style questions [18] (where lower is better) (Figure 8). Responses for the guided system were uniformly better, although no difference reached statistical significance.

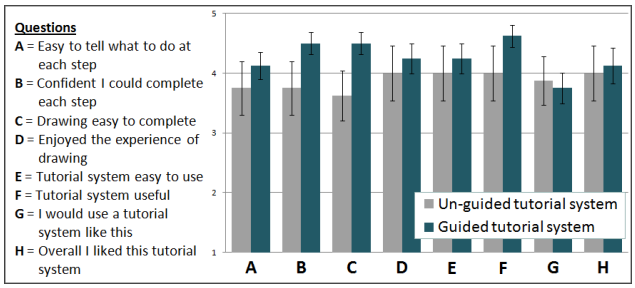
Questions
A = Easy to tell what to do at each step
B = Confident I could complete each step
C = Drawing easy to complete
D = Enjoyed the experience of drawing
E = Tutorial system easy to use
F = Tutorial system useful
G = I would use a tutorial system like this
H = Overall I liked this tutorial system
Un-guided tutorial system
Guided tutorial system
A B C D E F G H

Figure 7. Results from post-trial questionnaires (higher is better). Error bars show standard error.

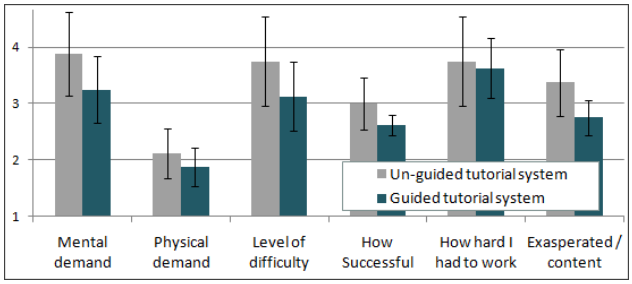

Figure 8. Post-trial questionnaire results from TLX-based questions (lower is better).

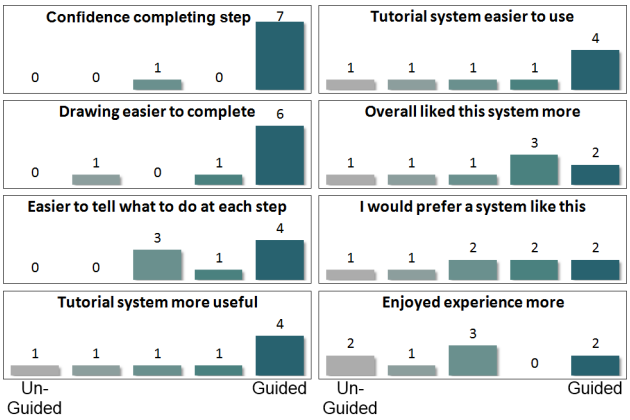
Confidence completing step
Tutorial system easier to use
Drawing easier to complete
Overall liked this system more
Easier to tell what to do at each step
I would prefer a system like this
Tutorial system more useful
Enjoyed experience more
Un-Guided
Guided

Figure 9. Subjective results comparing the un-guided and guided tutorial systems.

In the post-study questionnaire participants compared the two systems based on the first set of Likert questions. As can be seen in Figure 9, there is an indication that with the guided system, participants: were more confident they could complete each step; found the drawing easier to complete; and found it easier to tell what to do at each step. They also generally found the guided tutorial system more useful and easier to use versus the un-guided system. An interesting result is that users seemed to equally enjoy the experience of both the *guided* and *un-guided* condition. Our hypothesis is that while some users did find that the guided

system made things easier, they may have preferred a less guided and more free-form experience. To some extent, this validates our design goal of making the tutorial experience in Sketch-Sketch Revolution adaptable, as it allows users to configure the experience to have less guidance for users who prefer that type of experience.

Observations and User Feedback

Several interesting observations and comments were received from users, which we discuss below. Observations are grouped by the relevant design goal, where applicable.

Engagement (D1): In the un-guided tutorial, users were sometimes unsure of what to draw (P1: "What am I drawing?" P6: "Do I do the outline?") or precisely when they were finished (P1: "Is that all?" P5: "Next, I guess?"). This was more common in erasing and small steps.

Some nuances in steps were also missed in the un-guided tutorial. In particular, in one step in the fish drawing the author used the eraser to make a hard edge on grey shading created with the Airbrush. Three of the four users who drew the fish with the un-guided system did not notice the hard edge and instead used the eraser just to eliminate the excess grey shading outside of the border of the fish.

One participant looked at the author's supplementary text information to determine what to do when he was unsure. Most users repeatedly compared the author's step image on the secondary monitor to their own to determine the action they should carry out.

In the guided tutorial system, users typically knew what to do at each step (because the preview strokes played), and when they were finished (because the strokes for that step would finish). Sometimes they were unsure because they did not notice the preview stroke, and they would replay it.

Overall, users did appear to be engaged with the guided tutorials. For example, P1 began with the un-guided tutorial and (unprompted) verbalized his thought process throughout. During the guided tutorial, he hardly spoke. He was clearly engrossed in the drawing steps, previewing and replicating strokes in the exact manner as the author.

Application Learning (D4): Users who began their session with the un-guided tutorial system had difficulty interacting with the marking menus for layer manipulations, even with the positional callout instructions. In the guided system, where the callouts contained image instructions, three of the four users figured out how to use the marking menu immediately. The one user who did not (P5) confessed he did not "read the instructions" in the callout but, once he read them, he figured it out immediately.

One user (P2), who used the un-guided system first, did not understand how to use the Airbrush. She attempted to add colour to her drawing by tapping but was not getting the same results as in the author's image. It took her numerous attempts before understanding the interaction. Conversely, all users who started with the guided system used the Airbrush correctly since they could see in the preview stroke that the author dragged the pen across the screen.

Regardless of which condition users began with, all users appeared to be comfortable using SketchBook Pro towards the end of their session. Users performed complex interactions, such as layer manipulations and brush resizing, quickly and without assistance, suggesting that we achieved our design goal of providing Application Learning (*D4*).

In their questionnaire comments, 5 users explicitly stated that they thought the guided tutorial was helpful in teaching them how to use SketchBook Pro. (P2: "This tutorial actually had pop-ups describing where and how to use controls, which made it much better than the [un-guided]." P3: "I would use the guided first until I learned the system enough to practice in the un-guided setting." P5: "The guided tutorial moved me through the program quite quickly and has left me with a good base in a program that I was unfamiliar with prior to participating.")

Learning to draw (D5): Participants indicated an appreciation for the artistic guidance and assistance. (P2: "I think this tutorial is really good for beginners, since it shows you the exact moves to make and even auto-corrects. I like that you can control how much help it provides so that you can gradually become more independent.")

P6 began with the un-guided tutorial and had difficulty doing precision erasing. He moved the stylus rapidly back and forth in a scribbling motion, accidentally erasing edges of objects where precision was necessary. The participants who began with the guided tutorial system could see when drawing tasks required, slower, precise movements, and could mimic the speed of the previewed author strokes. So it did appear that users were learning certain drawing strategies from using our system (*D5*).

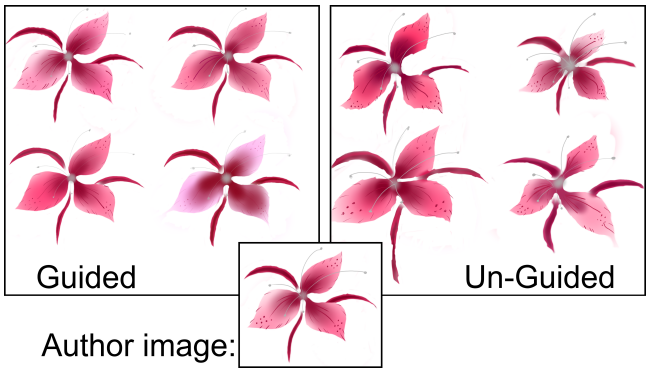

Figure 10. Participants' completed orchid drawing.

Success and Confidence (D6): P2 was not confident in her ability to recreate the drawing with the un-guided system, saying, "It's probably going to look crappy". P6 was also not confident, saying "Oh wow" when presented with a complicated drawing step. P8 sarcastically said "Are you sure this isn't April 1st?" at the start of the same step (drawing details on the orchid petals). These observations indicated that the un-guided tutorial could cause anxiety, possibly preventing a state of Flow (Figure 2).

There was a perceivable difference in the quality of the content generated in the two conditions. Figure 10 shows

the results of the orchid, demonstrating that Sketch-Sketch Revolution can help users experience successful content creation (*D6*).

Usage of Features: Most users relied heavily on the previews to see what they should draw next. Some users had to be reminded of the previewing hotkeys during the initial stage guided tutorial when it was clear they were trying to determine what to do in a step.

Almost all users kept their non-dominant hand on the keyboard as they carried out the guided tutorial for quick access to the hotkeys. Users became engrossed in steps that involved many small pencil strokes, resting one hand on the keyboard, the other on the tablet, and quickly alternating between previewing and sketching.

P1 became particularly engaged with the preview strokes. Even when the remainder of an erasing step was obvious, he only erased as much as the author did in any given preview stroke. A few users, on the other hand, got the idea for the erasing steps and erased all excess colour at once without watching each preview stroke.

Several users frequently checked the stroke feedback for each stroke when sketching with the pencil. One user would not progress to the next stroke unless he received a high distance score. Not everyone found the feedback useful. On the system feature questionnaire, P7 said "I can see the results, so 93% seems unnecessary."

Challenges Encountered: At the start of P7's second drawing, even after having added several layers in the guided tutorial, the first instruction said to add a layer and she confessed that she didn't know how. This demonstrates a potential issue with guided tutorials: that users become *too* reliant on guides and do not explicitly learn. This challenge has also been observed with video tutorials [25].

Although preview strokes were one of the most popular features among users, the previews sometimes demonstrate temporal dynamics that the user simply cannot replicate. For example, P6 saw how the author was able to erase very accurately in long ballistic strokes. P6 replicated the author's speed, but not his accuracy, suggesting that this particular author was not his ideal teacher.

A final challenge that appeared exclusively in the un-guided system was that errors made earlier in the tutorial could propagate and become more problematic later. For example, P1's drawing was much smaller than the artists and subsequently when he selected the eraser and set its size, it was too large for his drawing and he accidentally erased too much. In the un-guided tutorial, P4 missed one of the three leaves in the original author image. Later in the tutorial, when the leaves were being coloured, she saw that it was missing, remarking "Crap!". While not observed, this type of issue could also occur in the guided tutorial system, and we will discuss it in our future work section.

DISCUSSION

We have presented a new system, Sketch-Sketch Revolution, for automatically creating sketching tutorials that are *in-application*, *content-centric*, and *reactive.* Numerous novel features were implemented to satisfy a set of six design goals, such as providing an engaging and adaptable experience, supporting both application and domain learning, and ensuring the experience is successful.

The results of our qualitative user study were encouraging. Overall the most popular features of our system were the in-depth callouts, stroke assistance, and preview strokes. The in-depth callouts go beyond the "notes" used in Stencils-based tutorials as they provide detailed instructions and images for how to use widgets requiring more involved interactions. This seemed particularly effective for application learning. The stroke assistance was observed to reduce anxiety within users, and also resulted in greater success with the generation of content. The preview strokes were effective as they provided users with clear goals and also helped users learn subtle strategies about the art of drawing.

Our system was developed as an add-on for SketchBook Pro. It is interesting to think about how well our work applies to other application domains. While some specific techniques may not generalize, the higher level design goals which they adhere to should. For example, the stroke guidance algorithm may not be useful in other application domains, but similar application-dependent correction features could be implemented to *enable users to experience success* (D6). The only important distinction would be that the design goal *teach users to draw* would be modified to *teach users the domain*. However, we feel that our work is most relevant to software that requires involved interactions that are prevalent in design software. The actual implementation was facilitated by Sketchbook's existing journaling capabilities. While most of today's software systems have such capabilities, implementation may be more involved if there is no existing journaling mechanism.

FUTURE WORK AND CONCLUSION

There are a number of opportunities for future work.

First, an important follow-up evaluation would study the impact of our system on short- and long-term learning effects for both application and artistic learning. In particular, it would be interesting to formally measure a user's drawing abilities and expertise with SketchBook Pro after extended exposure to our system.

To improve application learning and prevent users from becoming too reliant on the instructional callouts, the amount of assistance could adapt to the user's behaviors and gradually reduce. For example, the system could recognize that a user has become adept at using marking menus, and automatically turn off those instructional callouts. Similarly, the system could detect when a user's sketching is accurate, and gradually reduce the stroke guidance value. The level of adaptation of our system could also be expanded by having the system intelligently choose tutorial content to match a user's skill set.

Our correction algorithm was limited to Pen and Pencil strokes. Some initial attempts to develop correction algo-

rithms for Airbrush and Paintbrush strokes were made, but it turned out to be fairly complex to parameterize and compare two brush patterns. In the future we aim to further investigate possible correction algorithms for additional brushes, such as the Airbrush, Paintbrush and Erasers.

Future implementations of the tutorial system could also contain more infrastructures to support creative freedom. While our system does allow choice between a spectrum of guided and un-guided experiences, the more they diverge from the author's content, the less relevant instructions and preview strokes may become. A future system version could, for example, detect when a user's drawing is out of proportion with the author's and make appropriate brush size suggestions. Regarding content creation, vision algorithms [14] could potentially be used to create preview strokes in the context of the user's drawing, even if it has taken on a much different shape than the author's.

In conclusion, we have created an authoring tool for sketching tutorials that records author actions and automatically generates interactive tutorials. The resulting tutorials were designed to provide an engaging experience for users, teach them how to use SketchBook Pro, assist and provide feedback on the content creation, and allow for customization in terms of assistance level and time commitment. A user study revealed our system to be useful and easy to use.

ACKNOWLEDGEMENTS

The authors thank Mark Davis for providing some of the initial ideas and insights that inspired this work, OCAD University and in particular Martha Ladly and Pedro Bonatto de Castro for their discussions, and Paul Telisman for generating tutorial content.

REFERENCES

1. Akers, D., Simpson, M., Jeffries, R. and Winograd, T. (2009). Undo and erase events as indicators of usability problems. *ACM CHI*. 659-668.
2. Arvo, J. and Novins, K. (2000). Fluid Sketches: Continuous Recognition and Morphing of Simple Hand-Drawn Shapes. *ACM UIST*. 73-80.
3. Bae, S.H., Balakrishnan, R. and Singh, K. (2008). ILoveSketch: As-natural-as-possible sketching system for creating 3D curve models. *ACM UIST*. 151-160.
4. Bae, S.H., Balakrishnan, R. and Singh, K. (2009). EverybodyLovesSketch: 3D Sketching for a Broader Audience. *ACM UIST*. 59-68.
5. Bederson, B.B. (2004). Interfaces for staying in the flow. *Ubiquity*. 5(27).
6. Carroll, J. M. (1990). *The Nurnberg Funnel*. MIT Press.
7. Carroll, J. M. and Rosson, M. B. (1987). *Paradox of the active user. Interfacing thought: cognitive aspects of human-computer interaction*. MIT Press. 80-111.
8. Chen, H., Wigand, R.T. and Nilan, M.S. (1999). Optimal experience of web activities. *Computers in human behavior*. 15(5), 585-608.
9. Csikszentmihalyi, M. (1988). The flow experience and its significance for human psychology. *Optimal Experience: Psychological Studies of Flow in Consciousness*. 15-35.
10. Farkas, D. K. (1993). The role of balloon help. *Journal of Computer Documentation*. 17(2):3-19.
11. Flagg, M. and Rehg, J.M. (2006). Projector-guided painting. *ACM UIST*. 235-244.
12. Fung, R., Lank, E., Terry, M. and Latulipe, C. (2008). Kinematic templates: end-user tools for content-relative cursor manipulations. *ACM UIST*. 47-56.
13. Google SketchUp Training. http://sketchup.google.com/intl/en/training/index.html
14. Grabler, F., Agrawala, M., Li, W., Dontcheva, M. and Igarashi, T. (2009). Generating photo manipulation tutorials by demonstration. *ACM SIGGRAPH*. 28(3):66.
15. Grossman, T., Fitzmaurice, G. and Attar, R. (2009). A survey of software learnability: metrics, methodologies and guidelines. *ACM CHI*. 649-658.
16. Grossman, T. and Fitzmaurice, G. (2010). ToolClips: An investigation of contextual video assistance for functionality understanding. *ACM CHI*. 1515-1524.
17. Grossman, T., Matejka, J. and Fitzmaurice, G. (2010). Chronicle: capture, exploration, and playback of document workflow histories. *ACM UIST*. 143-152.
18. Hart, S. and Staveland, L. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In P. Hancock & N. Meshkati (Eds.), Human mental workload. 139-183.
19. Huang, J. and Twidale, M.B. (2007). Graphstract: minimal graphical help for computers. *ACM UIST*. 203-212.
20. Igarashi, T., Matsuoka, S., Kawachiya, S. and Tanaka, H. (1997). Interactive beautification: a technique for rapid geometric design. *ACM UIST*. 105-114.
21. Kelleher, C. and Pausch, R. (2005). Stencils-based tutorials: design and evaluation. *ACM CHI*. 541-550.
22. Lee, Y.J., Zitnick, C.L. and Cohen, M. (2011). ShadowDraw: Real-Time User Guidance for Freehand Drawing. *ACM SIGGRAPH*.
23. Li, I., Nichols, J., Lau, T., Drews, C. and Cypher, A. (2010). Here's what I did: sharing and reusing web activity with ActionShot. *ACM CHI*. 723-732.
24. Nakamura, T. and Igarashi, T. (2008). An application-independent system for visualizing user operation history. *ACM UIST*. 23-32.
25. Palmiter, S. and Elkerton, J. (1991). An evaluation of animated demonstrations of learning computer-based tasks. *ACM CHI*. 257-263.
26. Singh, K. (1999). Interactive curve design using digital French curves. *ACM I3D*. 23-30.
27. Sukaviriya, P., Isaacs, E. and Bharat, K. (1992). Multimedia help: a prototype and an experiment. *ACM CHI*. 433-434.
28. Tuck, R. and Olsen, D. R. (1990). Help by guided tasks: utilizing UIMS knowledge. *ACM CHI*. 71-78.
29. Wobbrock, J.O., Wilson, A.D. and Li, Y. (2007). Gestures without libraries, toolkits or training: a $1 recognizer for user interface prototypes. *ACM UIST*. 159-168.